



Software

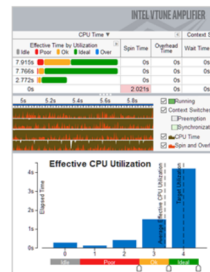
# Profiling your application with Intel® Vtune™ Amplifier

Paulius Velesko

# Getting a local copy of Vtune™

Get a free version of vtune for your PC/Mac/Linux machine:

<https://software.intel.com/en-us/vtune/choose-download>



## Only Intel® VTune™ Amplifier

This advanced profiler helps you increase application performance on modern hardware.

### Free Download

Download a free copy backed by community forum support.

[Download](#)

### Buy It Now

A paid product entitles you to:

- Private Priority Support from Intel's engineers
- The ability to submit confidential code samples
- Responsive help with technical questions
- Access to older versions
- Support renewals at a lower rate

Buy the software from a number of resellers or directly from the online store. Special pricing for academic research is available.

[Find a Reseller](#)

[Buy Now](#)

[Pricing for Academic Research](#)

#### Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Tuning at Multiple Hardware Levels

Exploiting all features of modern processors requires good use of the available resources

- Core
  - Vectorization is critical with 512bit FMA vector units (32 DP ops/cycle)
  - Targeting the current ISA is fundamental to fully exploit vectorization
- Socket
  - Using all cores in a processor requires parallelization (MPI, OMP, ... )
  - Up to 64 Physical cores and 256 logical processors per socket on Theta!
- Node
  - Minimize remote memory access (control memory affinity)
  - Minimize resource sharing (tune local memory access, disk IO and network traffic)

# Intel® Compiler Reports

FREE\* performance metrics

# Compile with -qopt-report=5

- Which loops were vectorized
  - Vector Length
  - Estimated Gain
  - Alignment
  - Scatter/Gather
- Prefetching
- Issues preventing vectorization
- Inline reports
- Interprocedural optimizations
- Register Spills/Fills

```
LOOP BEGIN at ../src/timestep.F(4835,13)
remark #15389: vectorization support: reference nbd(i) has unaligned access [ ../src/timestep.F(4836,16) ]
remark #15381: vectorization support: unaligned access used inside loop body
remark #15335: loop was not vectorized: vectorization possible but seems inefficient. Use vector always directive or -vec-threshold0 to override
remark #15329: vectorization support: irregularly indexed store was emulated for the variable <coefd(nbd(i))>, part of index is read from memory
remark #15305: vectorization support: vector length 2
remark #15399: vectorization support: unroll factor set to 4
remark #15309: vectorization support: normalized vectorization overhead 0.139
remark #15450: unmasked unaligned unit stride loads: 1
remark #15463: unmasked indexed (or scatter) stores: 1
remark #15475: --- begin vector cost summary ---
remark #15476: scalar cost: 4
remark #15477: vector cost: 4.500
remark #15478: estimated potential speedup: 0.880
remark #15488: --- end vector cost summary ---
remark #25439: unrolled with remainder by 2
LOOP END
```

# Getting your application ready for profiling

- Always add -g
  - No performance penalty
- On Cray Systems (Theta) add -dynamic
- E.x:
  - `cc -g -xMIC-AVX512 -dynamic -c test.cpp`

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



# Intel® Application Performance Snapshot

Profile your entire application at scale

# VTune™ Amplifier's Application Performance Snapshot

## High-level overview of application performance

- Identify primary optimization areas
- Recommend next steps in analysis
- Extremely easy to use
- Informative, actionable data in clean HTML report
- Detailed reports available via command line
- Low overhead, high scalability
- Can be run at scale - python ML/AI profiles

### Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.





# Usage on Theta

Launch all profiling jobs from **/projects** rather than **/home**

```
$ module load aps
```

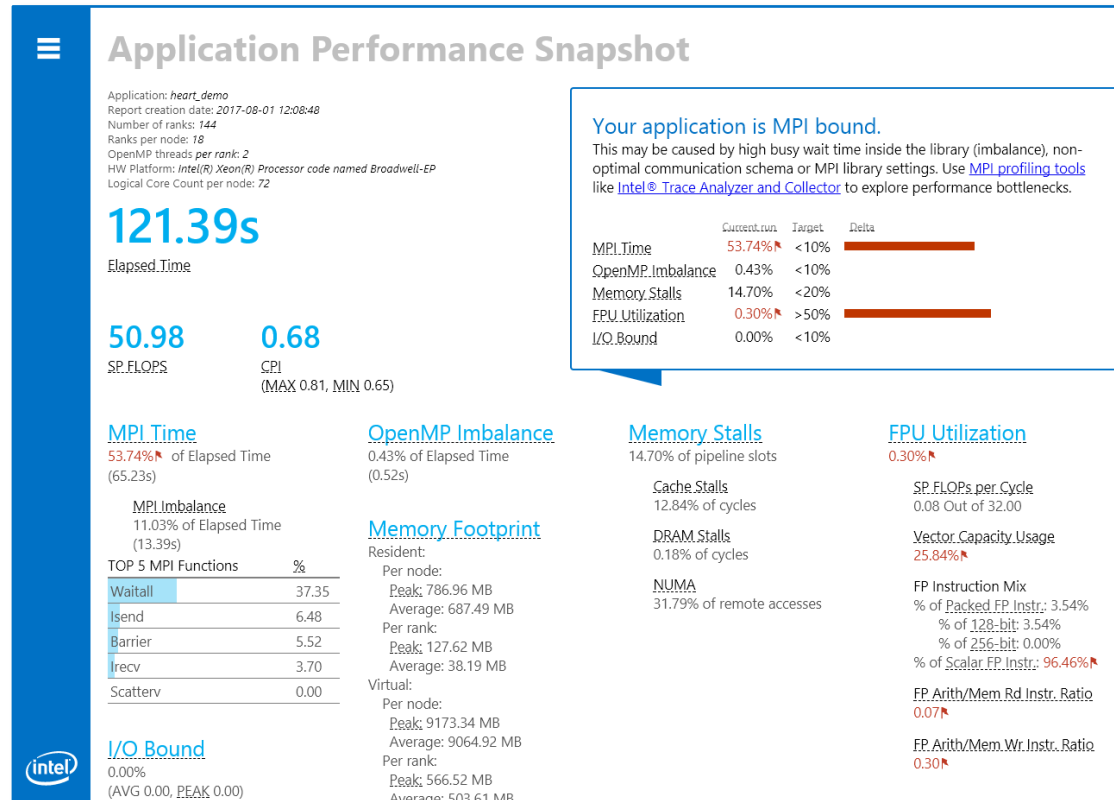
Launch your job in interactive or batch mode:

```
$ aprun -N <ppn> -n <totRanks> [affinity opts] aps ./exe
```

Produce text and html reports:

```
$ aps -report=./aps_result_ ...
```

# APS HTML Report



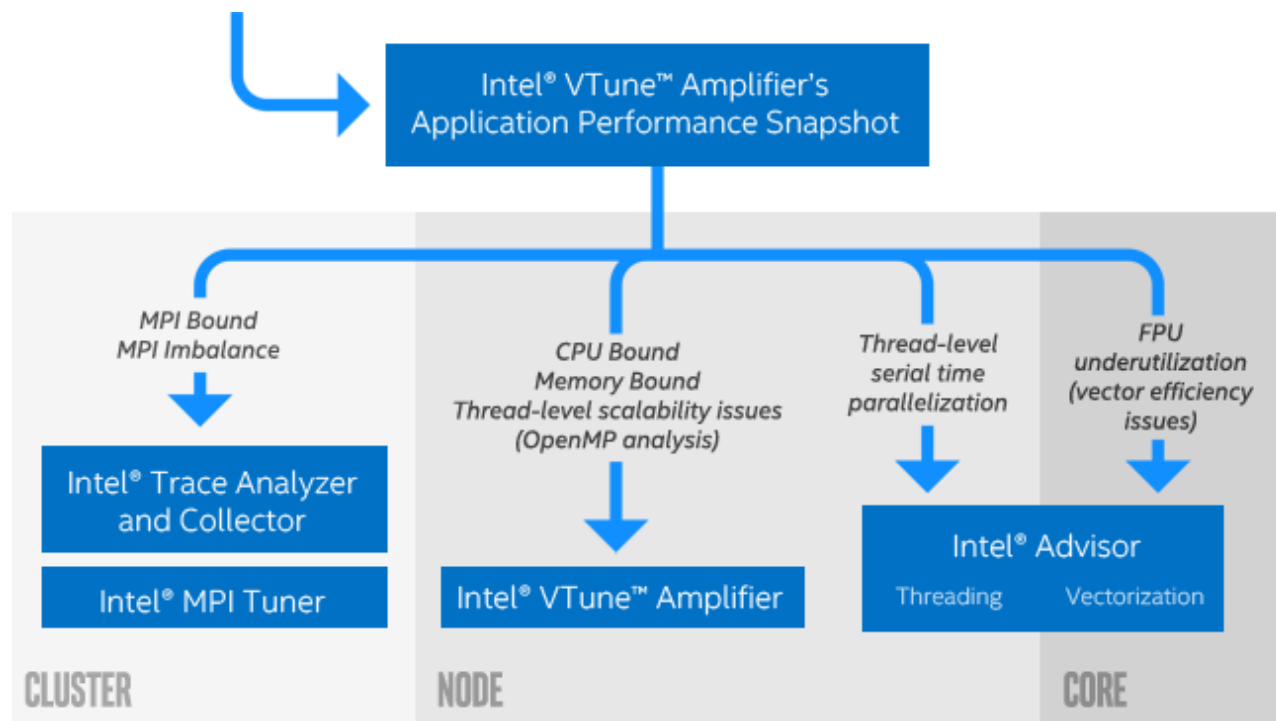
## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Tuning Workflow



## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Intel® VTUNE™ Amplifier

Core-level hardware metrics

<https://www.alcf.anl.gov/user-guides/vtune-xc40>

# Intel® VTune™ Amplifier

VTune Amplifier is a full system profiler

- Accurate
- Low overhead
- Comprehensive ( microarchitecture, memory, IO, treading, ... )
- Highly customizable interface
- Direct access to source code and assembly
- User-mode driverless sampling
- Event-based sampling

Analyzing code access to shared resources is critical to achieve good performance on multicore and manycore systems

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Predefined Collections

Many available analysis types:

- uarch-exploration General microarchitecture exploration
- hpc-performance HPC Performance Characterization
- memory-access Memory Access
- disk-io Disk Input and Output
- concurrency Concurrency
- gpu-hotspots GPU Hotspots
- gpu-profiling GPU In-kernel Profiling
- hotspots Basic Hotspots
- locksandwaits Locks and Waits
- memory-consumption Memory Consumption
- system-overview System Overview
- ...

Python Support

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Running on Theta

- Cray systems (such as Theta) use aprun instead of mpirun
  - No SPMD notation
    - `mpirun -n 1 amplxe-cl -c hotspots ./exe : -n <N-1> ./exe`
  - Use `$PE_RANK` in a bash script instead
    - If `$PE_RANK==0` `amplxe-cl -c hotspots ./exe`; else ...
  - `PMI_NO_FORK`
- Disable darshan
- Dynamic Linking

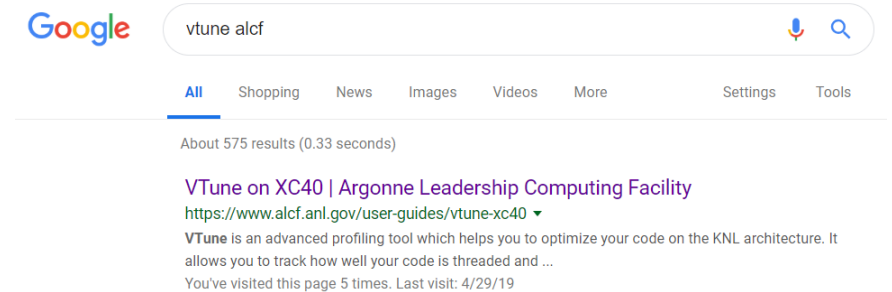
## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



# amplxe.qsub Script

- Copy and customize the script from `/soft/perftools/intel/vtune/amplxe.qsub`
- All-in-one script for profiling
  - Job size - ranks, threads, hyperthreads, affinity
  - **Attach to a single, multiple or all ranks**
  - Binary as arg#1, input as arg#2
    - `qsub amplxe.qsub ./your_exe arg1 arg2 ...`
  - Binary and source search directory locations
  - Timestamp + binary name as result directory
  - Save cobalt job files to result directory



## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.





# Viewing the result

- Text file reports:
  - `amplxe-cl -help report` How do I create a text report?
  - `amplxe-cl -help report hotspots` What can I change
  - `amplxe-cl -R hotspots -r ./res_dir -column=?` Which columns are available?
  - Ex: Report top 5% of loops, Total time and L2 Cache hit rates
  - `amplxe-cl -R hotspots -loops-only \`  
`-limit=5 -column="L2_CACHE_HIT, Time Self (%)"`
- Vtune GUI
  - `amplxe-gui`

Intel VTune Amplifier 2018

### Choose Analysis Type

Analysis Target | Analysis Type

- Algorithm Analysis
  - Basic Hotspots
  - Advanced Hotspots
  - Concurrency
  - Locks and Waits
  - Memory Consumption
- Compute-Intensive Application Analysis
  - HPC Performance Characterization**
- Microarchitecture Analysis
  - General Exploration
  - Memory Access
  - TSX Exploration
  - TSX Hotspots
  - SGX Hotspots
- Platform Analysis
  - CPU/GPU Concurrency
  - System Overview
  - GPU Hotspots
  - GPU In-kernel Profiling
  - Disk Input and Output
- Custom Analysis

## HPC Performance Characterization

Analyze important aspects of your application performance, including CPU utilization with additional details on OpenMP efficiency analysis, memory usage, and FPU utilization with vectorization information. For vectorization optimization data, such as trip counts, data dependencies, and memory access patterns, try Intel Advisor. It identifies the loops that will benefit the most from refined vectorization and gives tips for improvements. The HPC Performance Characterization analysis type is best used for analyzing intensive compute applications. Learn more (F1)

⚠ Vectorization analysis is limited for this platform. Only metrics based on binary static analysis such as vector instruction set will be available.

CPU sampling interval, ms

1

Copy Command Line to Clipboard@jlselgin2

Command line:

```
/soft/compilers/intel/vtune_amplifier_2018.1.0.535340/bin64/amplxe-cl -collect hpc-performance -app-working-dir /usr/bin -- ls
```

☐ Use -collect-with action

☒ Hide knobs with default values

Copy Close

Start

Start Paused

Choose Target

Command Line...

# Hotspots analysis for nbody demo

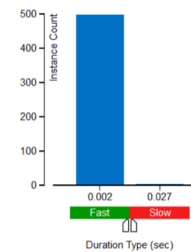
- `qsub amplxe.qsub ./your_exe ./inputs/inp`



## OpenMP Region Duration Histogram

This histogram shows the total number of region instances in your application executed with a specific duration. High number of slow instances may signal a performance bottleneck. Explore the data provided in the Bottom-up, Top-down Tree, and Timeline panes to identify code regions with the slow duration.

OpenMP Region: `startSomp$parallel.64@unknown.146.182`



Lots of spin time indicate issues with load balance and synchronization

Given the short OpenMP region duration it is likely we do not have sufficient work per thread

Let's look at the timeline for each thread to understand things better...

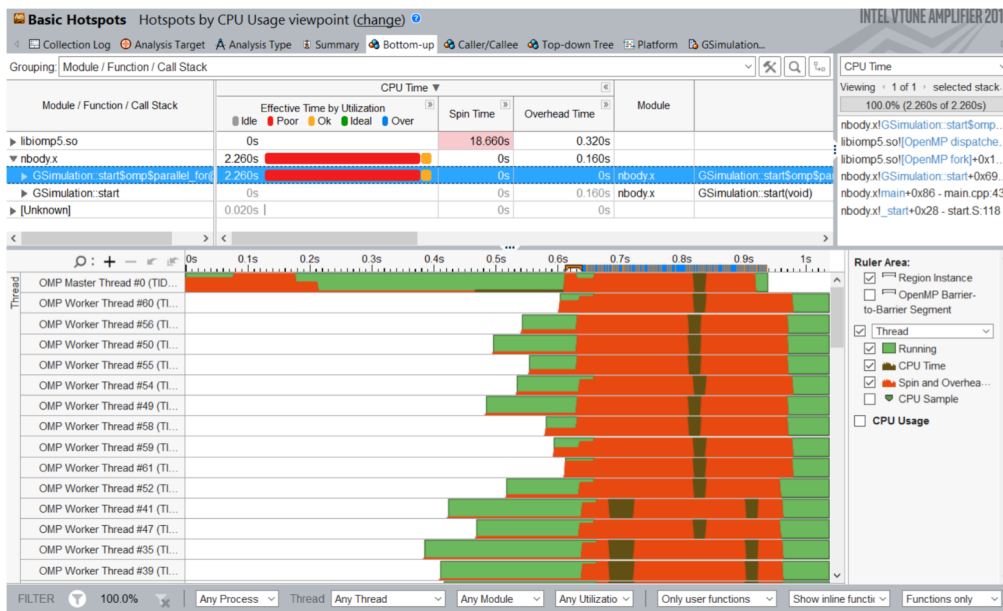
## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



## Bottom-up Hotspots view

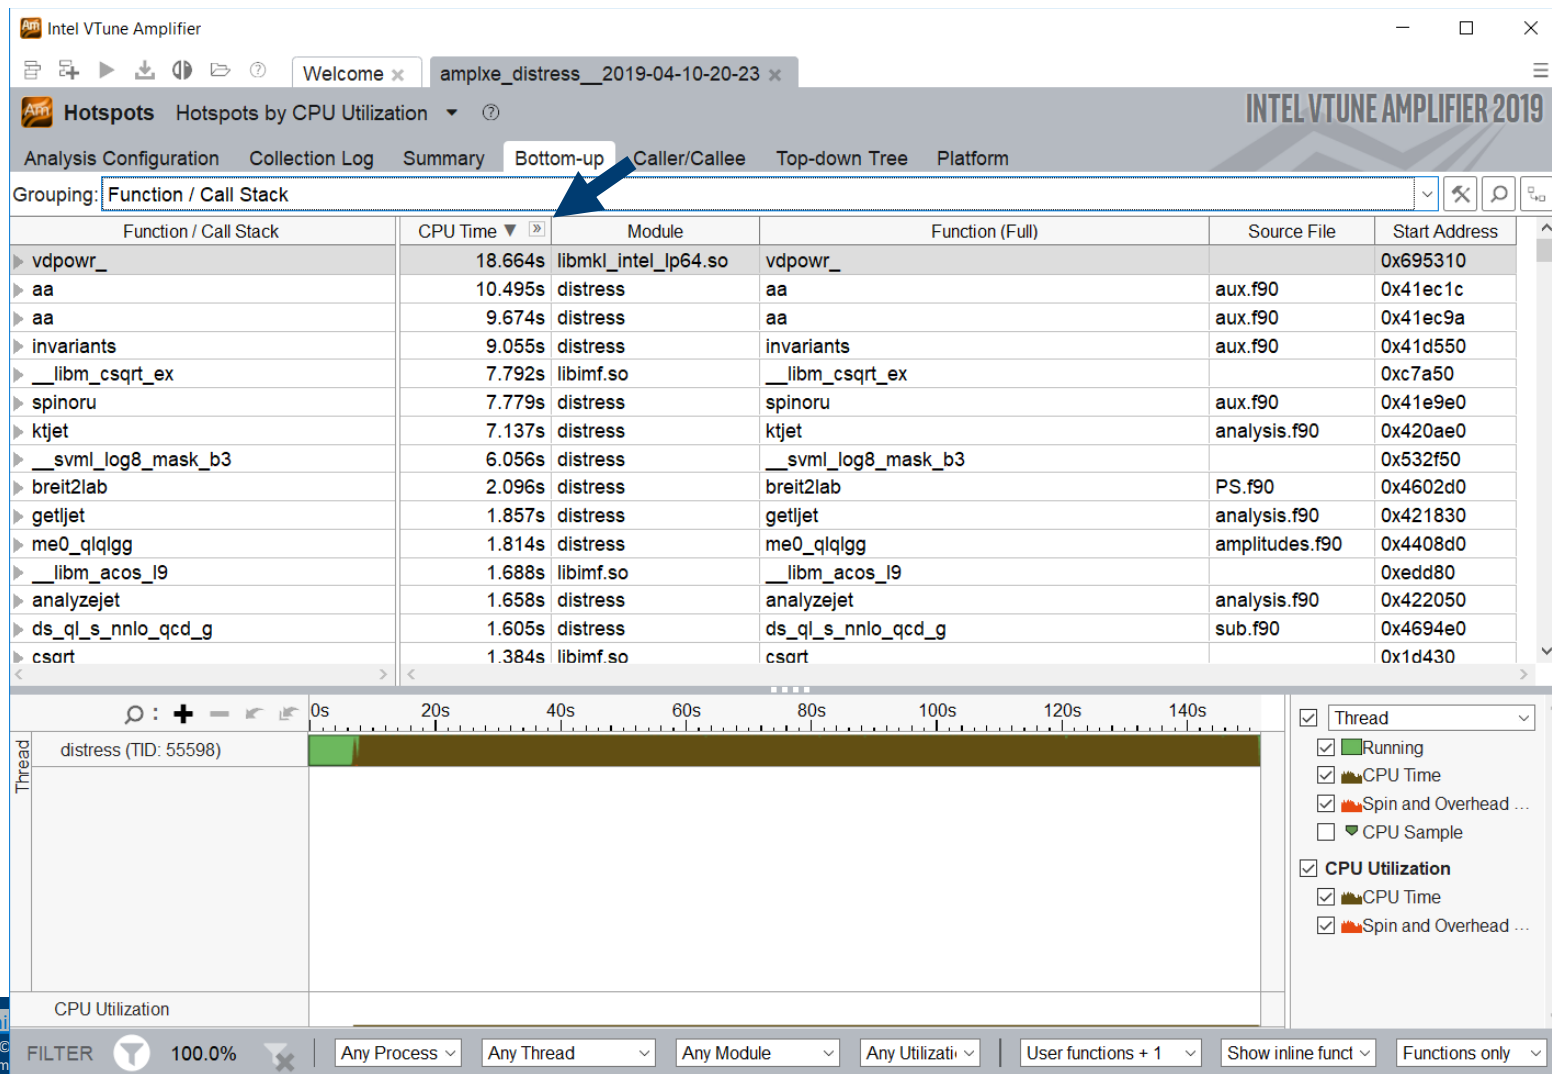


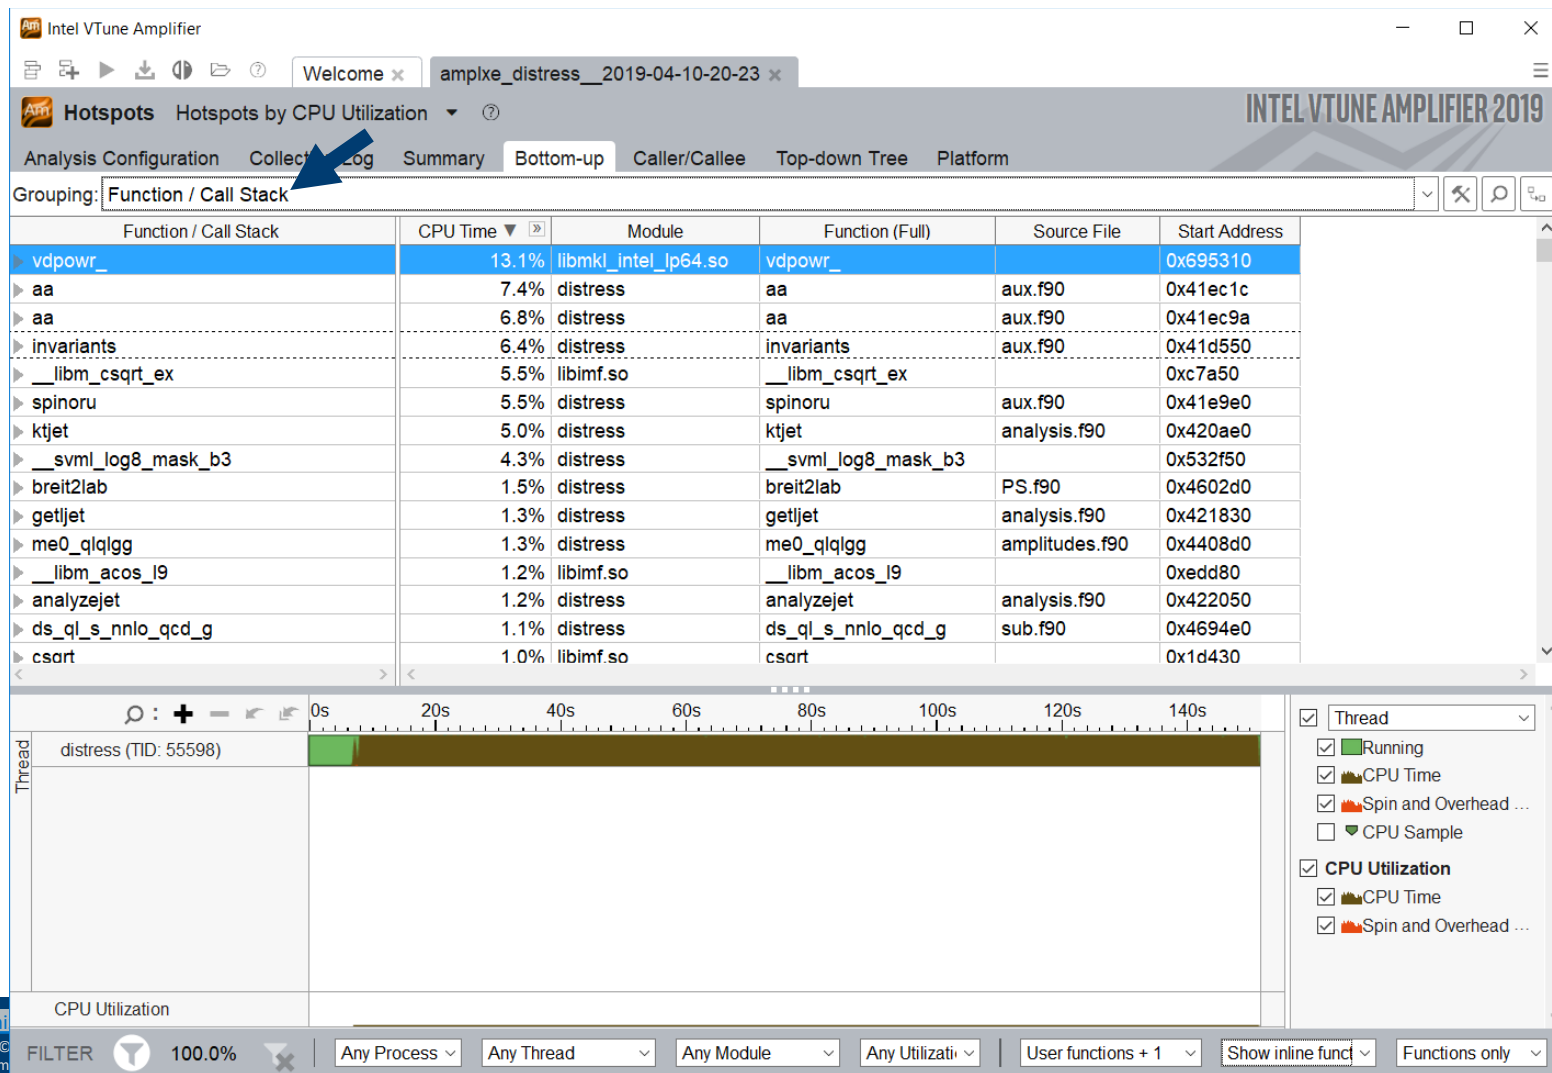
There is not enough work per thread in this particular example.

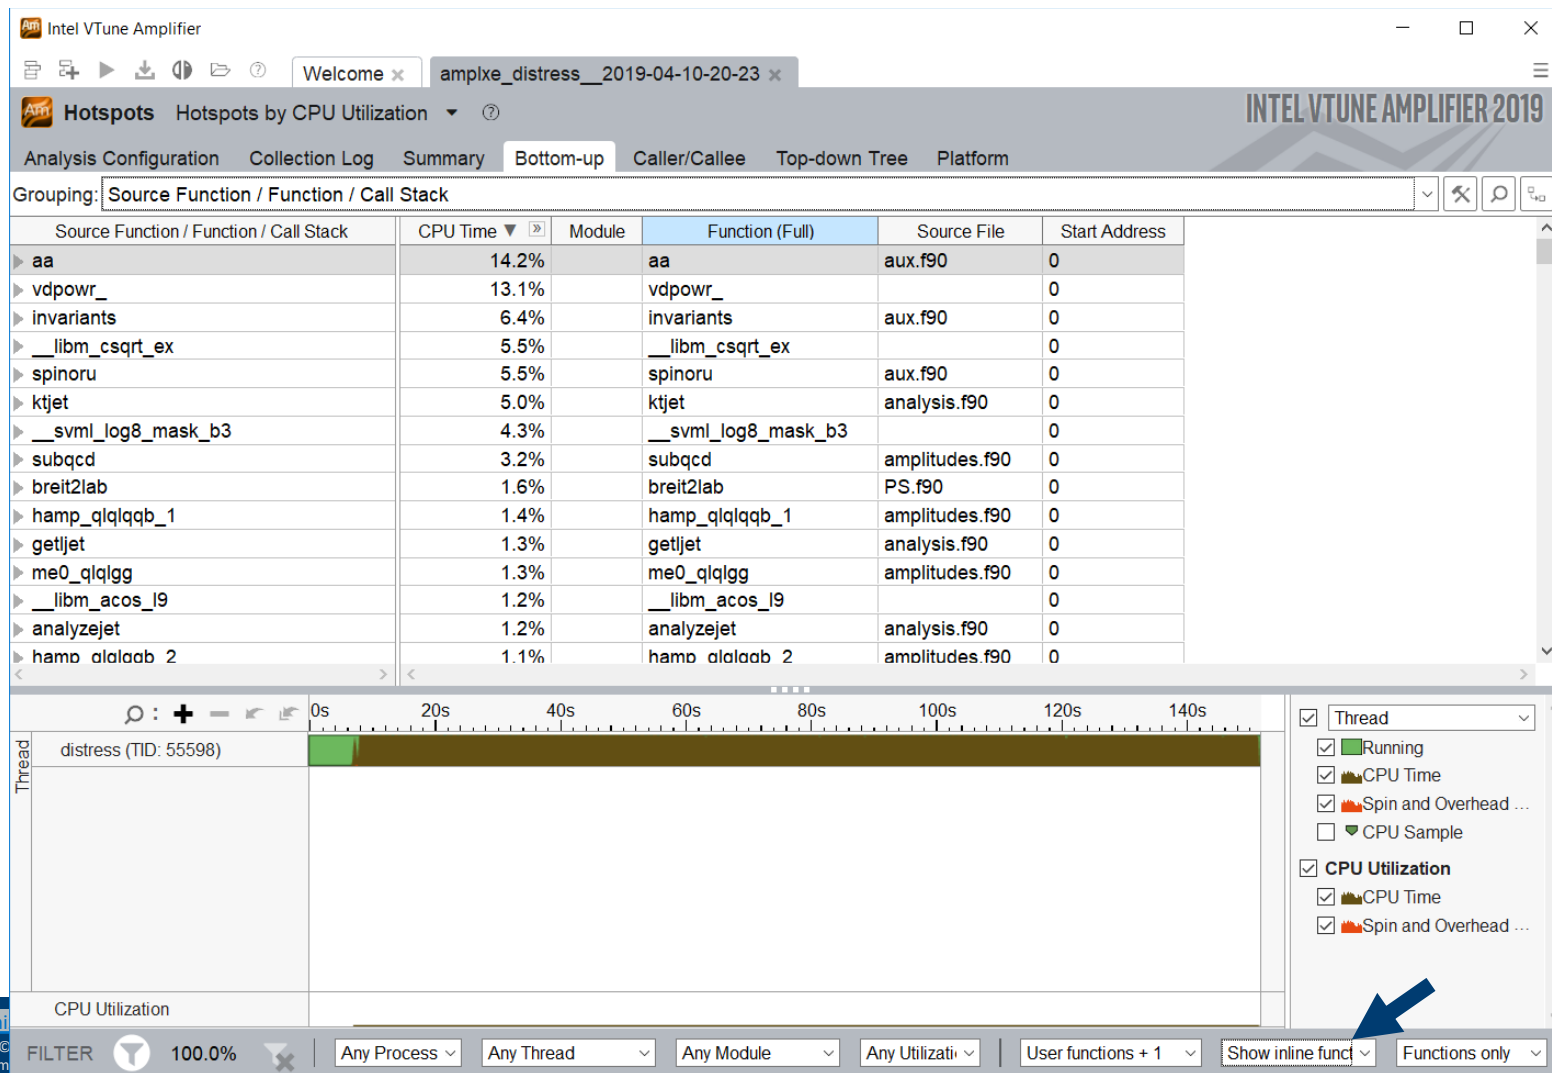
Double click on line to access source and assembly.

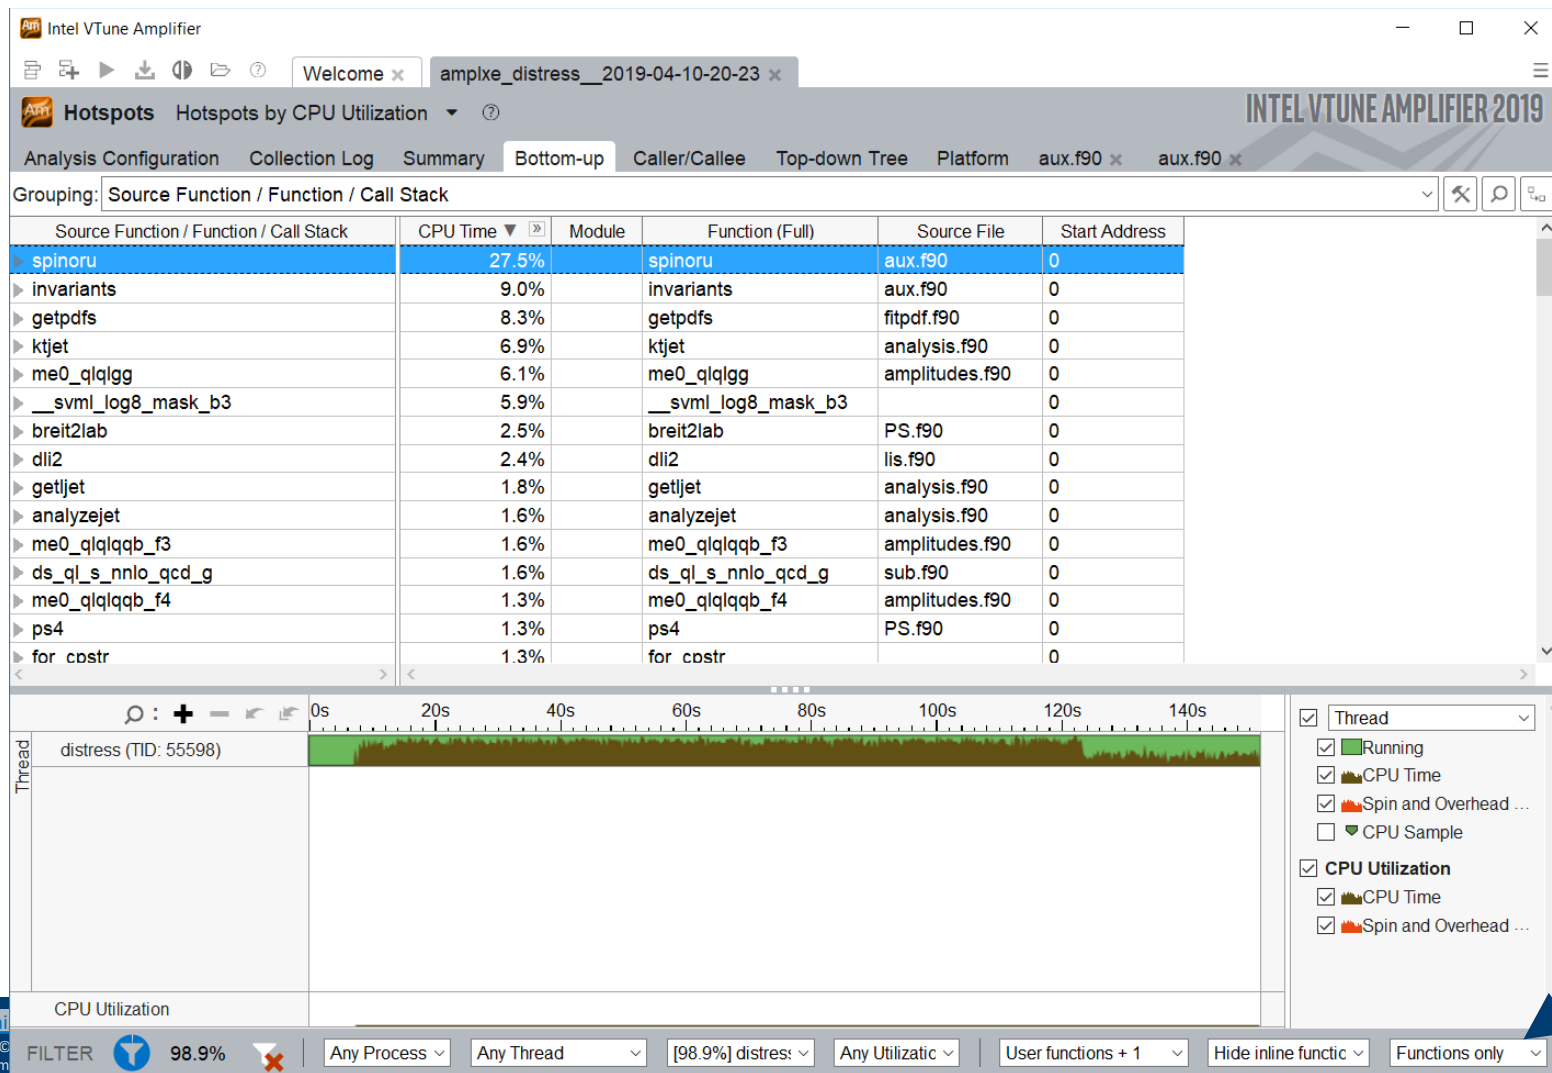
Notice the filtering options at the bottom, which allow customization of this view.

Next steps would include additional analysis to continue the optimization process.



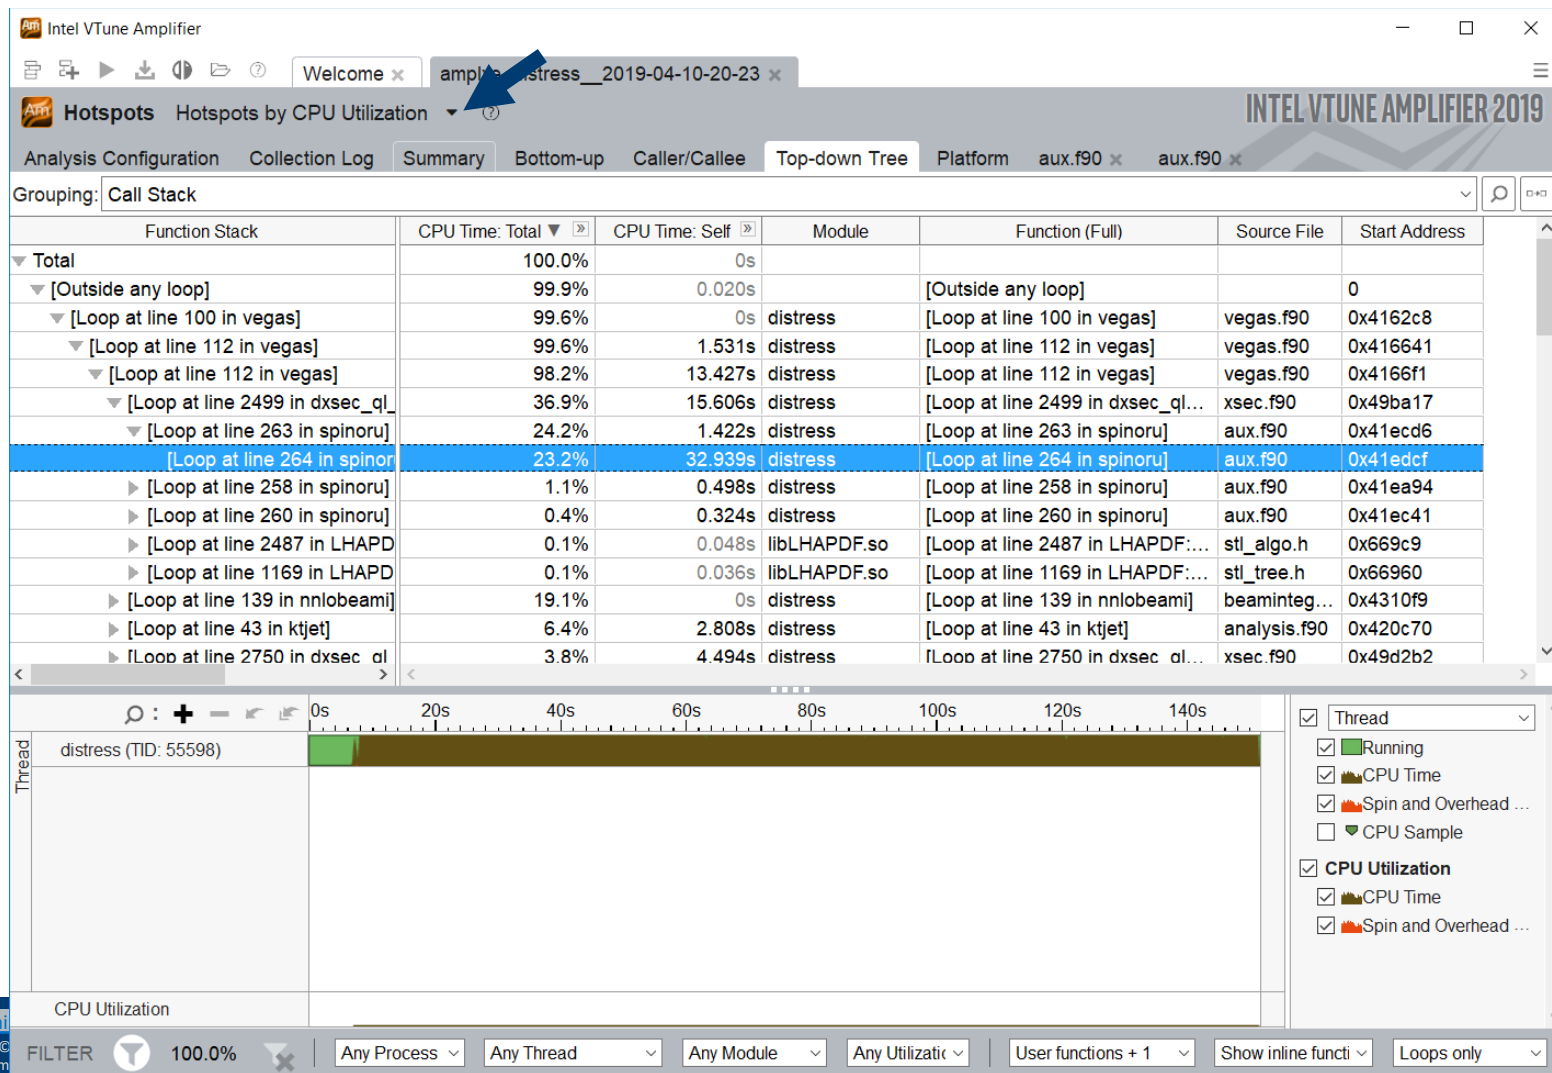


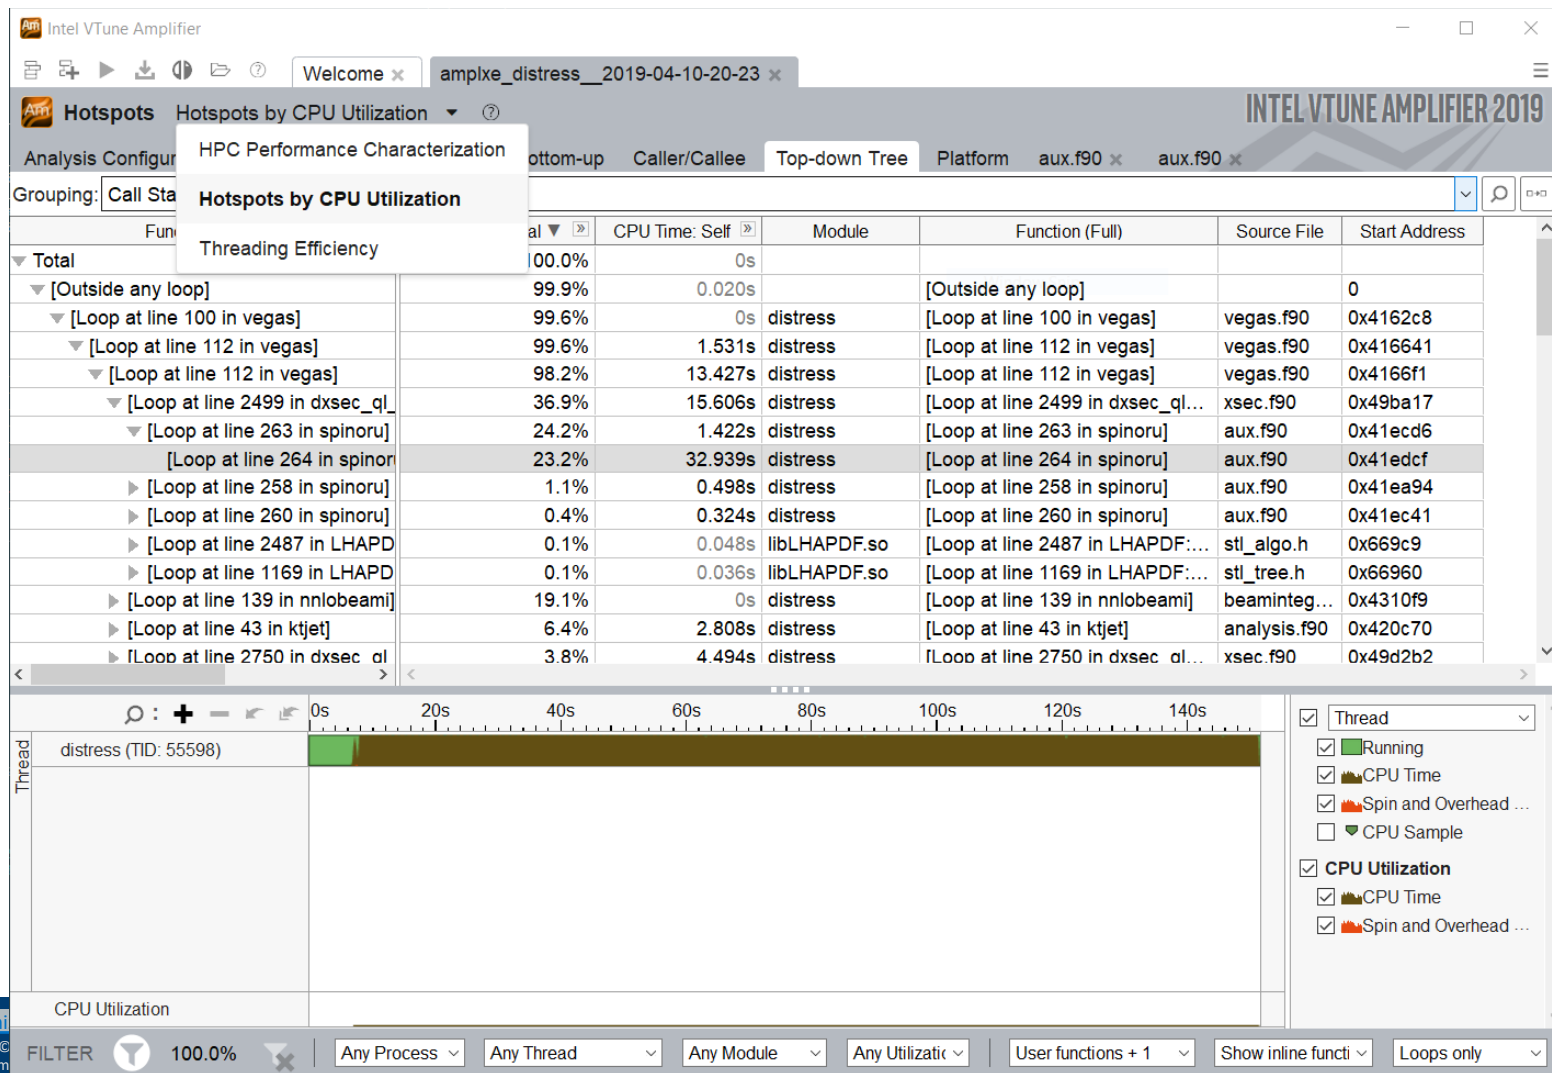






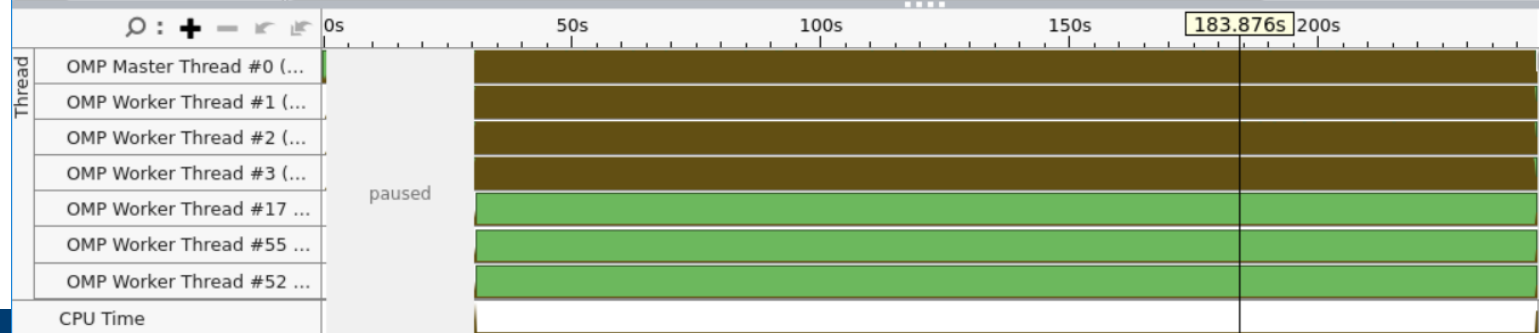






Grouping: Function / Call Stack

Function / Call Stack	CPU Time	CPI Rate	Front-End Bound	Bad Speculation	Back-End Bound					
					Memory Latency					Mer
					L1 Hit Rate	L2 Hit Rate	L2 Hit Bound	L2 Miss Bound	UTLB Overhead	
bicub_interpol1_aio_vec	26.8%	1.092	15.2%	2.3%	97.9%	100.0%	12.2%	0.0%	0.1%	0.0%
bicub_interpol2_aio_vec	11.1%	1.488	36.4%	0.9%	97.8%	100.0%	7.2%	0.0%	0.3%	0.0%
efield_gk_elec2_vec	10.9%	1.850	29.2%	1.0%	85.2%	100.0%	31.0%	0.0%	2.7%	0.0%
derivs_elec_vec	8.7%	2.241	57.9%	0.2%	86.2%	100.0%	28.7%	0.0%	0.3%	0.0%
field_following_pos2_vec	5.7%	0.969	43.6%	1.8%	94.3%	100.0%	33.3%	0.0%	0.2%	0.0%
i_interpol_ider0_aio_vec	5.3%	1.896	12.0%	0.0%	89.5%	100.0%	11.8%	0.0%	0.5%	0.0%
field_vec	4.8%	2.413	57.1%	0.0%	89.9%	100.0%	23.6%	0.0%	0.0%	0.0%
derivs_single_with_e_elec	3.0%	1.734	55.5%	0.0%	88.5%	100.0%	34.4%	0.0%	0.8%	0.0%
fld_vec_modulefield_follo	3.0%	1.189	34.9%	6.7%	74.0%	100.0%	73.0%	0.0%	0.9%	0.0%
bvec_interpol_vec	2.9%	1.131	38.8%	0.0%	91.2%	100.0%	36.2%	0.0%	0.0%	0.0%
pushe_single_vec	2.3%	1.943	43.9%	1.5%	71.3%	100.0%	54.7%	0.0%	1.1%	5.1%
i_interpol_ider0_aio_vec	1.8%	2.803	42.0%	0.1%	90.6%	0.0%	0.0%	0.0%	1.4%	0.0%



- Thread
- ☒ Running
  - ☒ CPU Time
  - ☒ CPU Time



# Profiling Python

# Python

Profiling Python is straightforward in VTune™ Amplifier, as long as one does the following:

- The “application” should be the full path to the python interpreter used
- The python code should be passed as “arguments” to the “application”

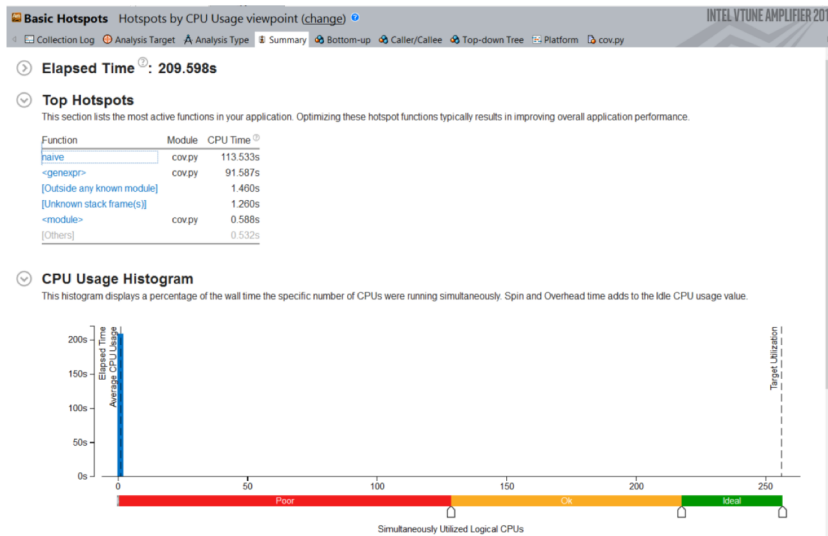
In Theta this would look like this:

```
aprun -n 1 -N 1 amplxe-cl -c hotspots -r res_dir \  
      -- /usr/bin/python3 mycode.py myarguments
```

Or use `amplxe.qsub` script as a starting point

# Simple Python Example on Theta

```
aprun -n 1 -N 1 amplxe-cl -c hotspots -r vt_pytest \  
-- /usr/bin/python ./cov.py naive 100 1000
```



Naïve implementation of the calculation of a covariance matrix

Summary shows:

- Single thread execution
- Top function is “naive”

Click on top function to go to Bottom-up view

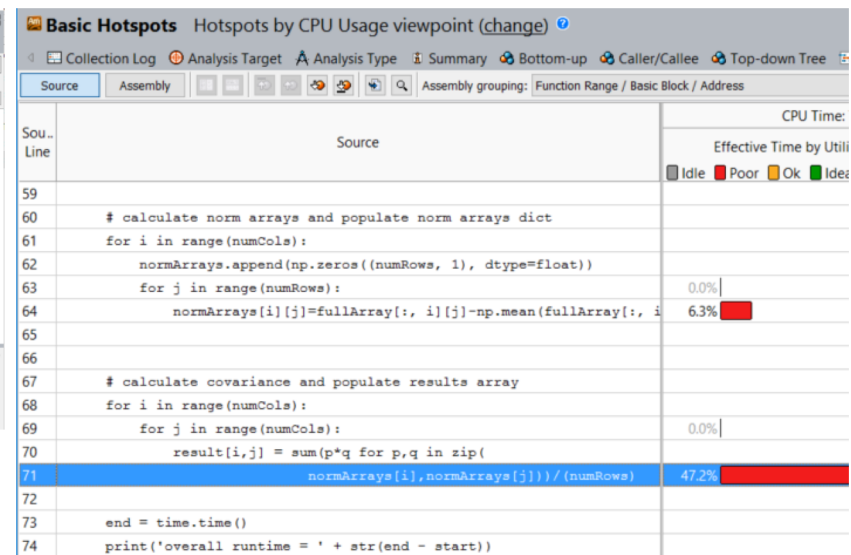
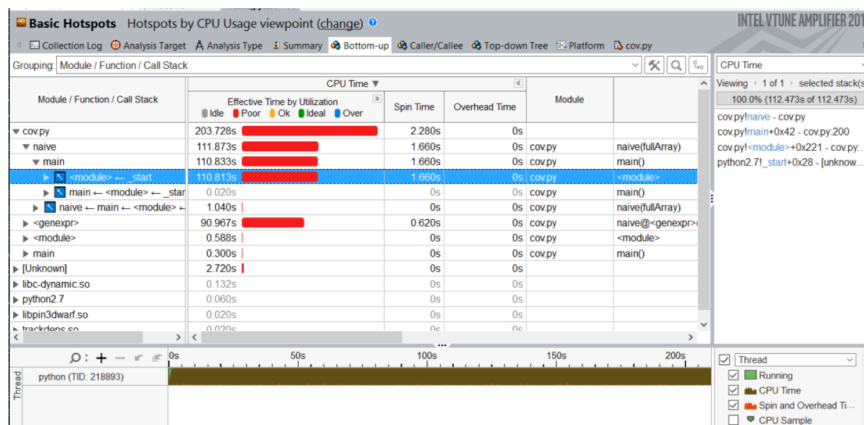
## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Bottom-up View and Source Code



Inefficient array multiplication found quickly  
We could use numpy to improve on this

Note that for mixed Python/C code a Top-Down view can often be helpful to drill down into the C kernels

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.





# Python & ML/DL

Vtune overheads scale with number of threads

First, ask yourself what is exactly that you want to find out?

- MPI call cost distribution can be acquired running aps at scale
- MKL call information will be given if you do `MKL_VERBOSE=1 python job.py`

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.





# Managing overheads

# ITT\_NOTIFY

- Instrument your code with `itt_pause()` and `itt_resume()`
  - Start, stop, detach and finalize your collection
  - Great when
    - Interested in specific region of code
    - Difficult to adjust overall runtime of your application
- Available in Fortran, C/C++, and Python
- Works with APS, Vtune, and Advisor

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# APS

- `aps` is meant to scale so has least amount of built-in knobs
  - Does work with `itt_notify`
- E.x: Profile the training portion of ML app, excluding setup and finalization
  - Insert `itt_resume()` before training step
  - Insert `itt_detach()` after training step
- `aprun aps --start-paused python ./app --arg1 --arg2`
  - Run paused, resume before training, detach and finalize after training
    - Job can be killed after `itt_detach()` is called

# Advisor

- Survey has to be run always
  - Overheads minimal
- Use mark-up-list tripcounts/map/dependencies only for regions of interest
  - `advixe-cl -R survey --project-dir ./res ...`
  - `advixe-cl -c tripcounts --mark-up-list=4 ...`
- `--interval=100`
- `--stop-after=1000` (1000 seconds)

→

ID	Function Call Sites and Loops	Total Time
4	[loop in GSimulation::start at GSimulation.cpp:208]	96.830s
7	[loop in GSimulation::start at GSimulation.cpp:248]	0.040s
1	[loop in fi_ini]	0.040s
2	[loop in fi_ini]	0.040s
3	[loop in GSimulation::init_vel at GSimulation.cpp:63]	0.010s
5	[loop in GSimulation::start at GSimulation.cpp:193]	96.830s
6	[loop in GSimulation::start at GSimulation.cpp:102]	96.900s

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Finalization

This step is inherently serial, takes a long time on KNL

SRC="--search-dir **src**:=/abs/path/to/source"

BIN="--search-dir **bin**:=/abs/path/to/binary"

- Advisor
  - aprun advixe-cl -c survey --project-dir test -no-auto-finalize
  - advixe-cl -R survey --project-dir test \$SRC \$BIN
- Vtune
  - aprun amplxe-cl -c hotspots -r test -no-auto-finalize
  - amplxe-cl -R summary -r test \$SRC \$BIN

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Common Issues

- I don't see my functions
  - Debug symbols
    - Do a text report, look for "Cannot find" warnings
    - Add those to --search-dir
    - Re-finalize

# Common Issues

- My python job hangs
  - Detach collection before your job finishes (--duration, --stop-after, or itt\_detach)
  - Reduce number of threads

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.





# Remember

Compile with -g and -dynamic

Profile 1 rank and small number of threads - `amplxe.qsub/advixe.qsub`

Advisor for big picture

Vtune for details

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Resources

## Product Pages

- <https://software.intel.com/sites/products/snapshots/application-snapshot>
- <https://software.intel.com/en-us/advisor>
- <https://software.intel.com/en-us/intel-vtune-amplifier-xe>

## Detailed Articles

- <https://software.intel.com/en-us/articles/intel-advisor-on-cray-systems>
- <https://software.intel.com/en-us/articles/using-intel-advisor-and-vtune-amplifier-with-mpi>
- <https://software.intel.com/en-us/articles/profiling-python-with-intel-vtune-amplifier-a-covariance-demonstration>

### Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.



# Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

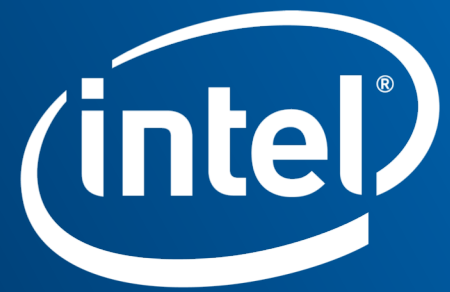
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Copyright © 2018, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

## Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804



Software

# When do I use Vtune vs Advisor?

## Vtune

- What's my cache hit ratio?
- Which loop/function is consuming most time overall? (bottom-up)
- Am I stalling often? IPC?
- Am I keeping all the threads busy?
- Am I hitting remote NUMA?
- When do I maximize my BW?

## Advisor

- Which vector ISA am I using?
- Flow of execution (callstacks)
- What is my vectorization efficiency?
- Can I safely force vectorization?
- Inlining? Data type conversions?
- Roofline

### Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.



# VTune Cheat Sheet

Compile with `-g -dynamic`

`amplxe-cl -c hpc-performance -flags -- ./executable`

- `--result-dir=./vtune_output_dir`
- `--search-dir src:=../src --search-dir bin:=./`
- `-knob enable-stack-collection=true -knob collect-memory-bandwidth=false`
- `-knob analyze-openmp=true`
- `-finalization-mode=deferred` if finalization is taking too long on KNL
- `-data-limit=125` ← in mb
- `-trace-mpi` for MPI metrics on Theta
- `amplxe-cl -help collect survey`

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.  
\*Other names and brands may be claimed as the property of others.

<https://software.intel.com/en-us/vtune-amplifier-help-amplxe-cl-command-syntax>



# Advisor Cheat Sheet

Compile with `-g -dynamic`

`advixe-cl -c roofline/dependencies/map -flags -- ./executable`

- `--project-dir=./advixe_output_dir`
- `--search-dir src:=../src --search-dir bin:=./`
- `-no-auto-finalize` if finalization is taking too long on KNL
- `--interval 1` (sample at 1ms interval, helps for profiling short runs)
- `-data-limit=125` ← in mb
- `advixe-cl -help`

# nbody - ver5\_mpi

```
[pvelesko@jlsellogin1 ver5_mpi]$ mpirun -n 2 advixe-cl -c survey --project-dir SDL -- ./nbody.x
Intel(R) Advisor Command Line Tool
Copyright (C) 2009-2019 Intel Corporation. All rights reserved.
Intel(R) Advisor Command Line Tool
Copyright (C) 2009-2019 Intel Corporation. All rights reserved.

advixe: Collection started.
advixe: Collection started.
=====
Initialize Gravity Simulation
nPart = 2000; nSteps = 500; dt = 0.1
-----
s      dt      kenergy      time (s)      GFlops
-----
50      5      0.1432      1.0434      5.5606
100     10      2.4341      1.0409      5.5738
150     15      8.1256      1.0419      5.5685
200     20     17.877      1.0423      5.5665
250     25     32.966      1.0401      5.578
300     30     55.786      1.0409      5.5738
350     35     91.132      1.0408      5.5747
400     40     150.12      1.0428      5.5638
450     45     264.78      1.0413      5.572
500     50     571.53      1.0405      5.5762

# Number Ranks      : 2
# Number Threads    : 1
# Total Time (s)     : 10.415
# Average Performance : 5.5717 +/- 0.0046592
=====
advixe: Collection stopped.
advixe: Collection stopped.
advixe: Opening result 21 % Resolving information for 'libc.so.6'
advixe: Warning: Cannot locate debugging information for file '/lib64/libc.so.6'.
advixe: Warning: Cannot locate debugging information for file '/lib64/libc.so.6'.
advixe: Opening result 100 % done
advixe: Preparing frequently used data 0 % done
advixe: Preparing frequently used data 100 % done

Elapsed Time: 10.43s
Total CPU time: 10.42

advixe: Opening result 99 % done
advixe: Preparing frequently used data 0 % done
advixe: Preparing frequently used data 100 % done

Elapsed Time: 10.43s
Total CPU time: 10.42

[pvelesko@jlsellogin1 ver5_mpi]$
[pvelesko@jlsellogin1 ver5_mpi]$
```

## Optimization Notice

Copyright © 2018, Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.

